ZLW

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

IN RE APPLICATION OF: Andrew Zador

CASE: 057688.010002

COMMUNICATION

SERIAL NO.: 10/730,189

FILED ON: December 5, 2003

FOR: Imaging System Utilizing Spatial Image Oscillation

ATTENTION OF
Examiner Richard M. Bemben
Art Unit 2622

COMMISSIONER OF PATENTS
WASHINGTON, D.C. 20231

Dear Sir:

If any charges or fees must be paid in connection with the following communication, they may be paid out of our Deposit Account No. 502428.

In accordance with the Rules, applicant is filing herewith a certified copy of the priority Canadian Application No. CA 2349828 for insertion into the file for the above-identified patent application, in support of Applicant's claim of priority thereto.

Acknowledgement of receipt of the certified priority document, and confirmation of Applicant's priority claim, are respectfully solicited.

GREENBERG TRAURIG LLP

77 WEST WACKER DRIVE
SUITE 2500
CHICAGO, ILLINOIS 60601
(312) 456-8400
CUSTOMER NO. 34018

Douglas B. Teaney	33459	Mark R. Galis	36805
Richard D. Harris	27898	Jeffrey P. Dunning	50686
Howard E. Silverman	32492	Jeffrey G. Mote	48791
Herbert H. Finn	38139	Cameron M. Nelson	55486
Gary R. Jarosik	35906	Kevin J. O'Shea	51041
James J. Lukas	59114	Eric Maiers	59614

Should anything further be required, a telephone call to the undersigned at (312) 456-8400 is respectfully requested.

Respectfully submitted,

Dated: 8\20\08

Douglas B. Teaney
One of attorneys for Applicant

CERTIFICATE OF MAILING

I hereby certify that this COMMUNICATION is being deposited with the United States Postal Service as First Class Mail under 37 C.F.R. §1.8, postage prepaid, in an envelope addressed to: Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450 on

Douglas B. Teaney

57,454,330

Office de la propriété intellectuelle du Canada

Un organisme d'Industrie Canada

Canadian Intellectual Property Office

An Agency of Industry Canada

Bureau canadien des brevets

Certification

La présente atteste que les documents ci-joints, dont la liste figure ci-dessous, sont des copies authentiques des documents déposés au Bureau des brevets.

Canadian Patent Office

Certification

This is to certify that the documents attached hereto and identified below are true copies of the documents on file in the Patent Office.

Specification and Drawings, as originally filed, with Application for Patent Serial No: **CA 2349828**, on June 6, 2001, by **ANDREW M. ZADOR**, for "Method, Apparatus, and System for Extracting Denoised, High-Resolution Images, Texture, and Velocity from a Lower-Resolution, Noisy Detector".

Agent certificateur/Certifying Officer

May 30, 2008

Date

Canada

(CIPO 68)
31-03-04

OPIC CIPO

Title:

METHOD, APPARATUS, AND SYSTEM FOR EXTRACTING DENOISED, HIGH-RESOLUTION IMAGES, TEXTURE, AND VELOCITY FROM A LOWER-RESOLUTION, NOISY DETECTOR.

Abstract:

A method, apparatus, and system for extracting a high-resolution, maximally detector-noise-removed image, along with texture, motion and velocity information, all directly from a detector array and its supporting logic, is described. The texture isolation, edge-motion detection, and noise removal processes rely largely upon principles deduced from mammalian vision. Spatially coincident opponent center/surround structures fed by the primary detector array are employed to remove spatio-temporally random detector-noise events. Static edges are accurately located by spatially oscillating the image with respect to the detector, filtering only for those edges whose motions reflect purely the induced oscillation, and obtaining accurate phase information at edge-crossings. Spatio-temporal activity in the image itself, not suitably matching the space-time characteristics of this oscillation (or a Doppler shifted version of it) can also be treated as noise if desired. During static periods (i.e. between induced oscillations), objects undergoing real motion detectable in the plane of the array can be selectively given attention by temporarily suppressing (either actively, or passively) those detectors having detected stationary edges during the previous oscillation, greatly reducing the computational load on robotic vision systems. Using Doppler techniques, accurate object velocities may be collected during oscillation, by detecting the frequency shift (with respect to the induced oscillation), of objects crossing multiple opponent center/surround detectors, and noting the orientation of the sequence of crossings. For a particular orientation of oscillation, center/surround crossings by a slice of texture in that orientation will yield a specific frequency spectrum of edge-crossing bursts. Given more than one orientation (and possibly scale) of such oscillation, and grouping regions by their resulting oscillation spectra, texture regions may be efficiently isolated and identified. Textures, edges and isolated points would have substantially different spectra, useful for identification. Oscillation of scale (zooming in and out) would permit radial velocity detection for simple collision avoidance detectors. A spatially-varying detection threshold, based upon localized image content, can be employed to suppress texture, to control entropy, or to match human vision characteristics. This detector does not employ raster scanning, but is instead event-driven, reporting exact timings and orientations of edge crossings from one opponent center/surround to the next; therefore, its spatial- and velocity-reporting accuracy are functions of its opponent center/surround temporal sensitivity, leaving open the possibility of using cheaper, larger pixels, with fewer pixels per detector array. Since actual velocities are reported in real-time from the array support logic, there is no need to calculate processor-intensive, inaccurate motion-vector fields (requiring a minimum of one frame of latency). Another benefit is the self-calibration of all adjacent detectors by systematic detector-difference cancellation during oscillation of the array, referenced against either a black perimeter (for absolute calibration), or against the average illumination over the entire field (which would permit color-constancy as in human vision).

Field Of The Invention:

The present invention relates to a method, apparatus, and system for the design of an image detector architecture to obtain a high-resolution, de-noised image, its texture-regions, and velocity-field, directly from a non-raster scanned, lower-resolution detector, removing spatio-temporal detector noise and pixel to pixel differences, while greatly reducing computational load for intelligent vision applications.

Brief Description Of The Drawings:

Preferred embodiments of the present invention will now be described, by way of example only, with reference to the attached Figures, wherein:

Figure 1 shows various proposed illuminated opponent center/surround structures, with output responses.

Figure 2 shows several primary detectors connected to a secondary opponent center/surround cell, interleaved with other primary detectors connected to another secondary opponent center/surround cell of opposite polarity.

Figure 3 shows several orientations of linear secondary opponent center/surround cells contributing to a resulting circular secondary opponent center/surround cell.

Figure 4 shows a graded neutral density wedge covering the edge of a detector array, to keep individual center/surround responses within the secondary detector's range, given uniform illumination, to allow greater total dynamic range for the array, and to give a reference-black level surround to the array.

Figure 5 shows spatially coincident opposite-polarity center/surrounds, and their role in canceling detector noise.

Figure 6 shows a single orientation of oscillation along a path sampling aligned slices of different textures and features, and their resulting one-dimensional center/surround crossing spectra.

Background Of The Invention:

In the field of radar the concept of the "chirp" is well known as a means of improving signal-to-noise ratio (SNR) electronically, by sweeping through frequencies in a controlled short signal burst and listening for a matching echo. It is used to cut through "radar clutter" and to circumvent electronic counter-measures (also known as "ECM" or "jamming"). This is accomplished by using a specific time-varying frequency and filtering any returned radio-echo, listening for the same time-varying signal or its Doppler-shifted echo. Noise, whether natural or artificial as in jamming, is highly unlikely to possess the same time-frequency chirp (which may be pseudo-randomly varied), and noise can thus be suppressed from incoming signals.

This first idea, thus far applied to modulating transmissions in the radio or acoustic signal domain, is proposed to be applied to spatially oscillating an image detector array (or identically, oscillating the image falling upon the detector array) in one or more dimensions in a controlled fashion. The detector array (or the image arriving upon it) is spatially oscillated in a known or determinable manner, and the received moving image data falling upon the array are filtered for the expected resulting motions. This action is believed by the author to have an analogue in human vision in the form of micro-saccadic motions, which, so far, have been assumed to exist exclusively to refresh the image information falling upon the retina (which ceases to transmit an image in the absence of such image-changing motions). This motion signature could be used to determine velocities of image components upon the detector array by spatial-

Doppler means, and to reduce the effects of time-varying detector or image noise. Such a "threshold extension" technique used to reject noise, increasing SNR, could be applied to radiology, possibly to permit reduction of the radiation dose required to produce an x-ray, and/or to bring diagnosable features up out of the signal noise. A hairline fracture oscillated on a detector array would have a constant shape, unlike the time-varying noise on top of it which might mimic hairlines in any radiological still-image.

In human vision, an image falling onto the cones of the retina is interpreted as a differential image by neural structures within the retina, which convert the image into difference signals between centers and their annular opponent-sense surrounds. These concentric structures come in two polarities; on-center – off-surround, and off-center – on-surround, with two (and possibly three) color pair variants; red-green opponents, blue-yellow opponents, (and possibly a synthesized black-white opponent). Other structures in the human visual system (ocular dominance columns) appear to indicate that these concentric opponent center/surround structures are based upon various orientations of similarly linearly opposed center-surrounds.

For the proposed monochrome detector array (easily extended to color, and to non-light applications such as ultrasound etc.), a linear, symmetric "secondary" opponent center/surround detector (fed by at least three "primary" luminance detectors) would, upon exposure to illumination of the primary detectors feeding it, output a fast-rise-time pulse followed by an analog level proportional to either the luminance difference or the log of the luminance difference (depending upon the application), between the center of the opponent center/surround detector and its surround. This level would have a characteristic polarity depending upon whether the opponent detector was on-center or off-center (the surround would be opponent to either variant of center), and the level would have a time-constant, decaying after initial exposure. It would also possess a turn-on threshold, which could be interpreted as a "Just Noticeable Difference" (JND) between the center and its surround. This threshold could be fixed, or tuned in real-time to imitate human sensitivity. Multiple opponent center/surround detectors of the same polarity chained in a specific orientation, and exposed to a moving luminance edge, would output a train of pulses (directly related to the apparent velocity of the edge) superimposed on the contrast level signal, representative of the exact times of successive center/surround crossings. Groups of these could be further chained in a multi-resolution structure, with different orientations separately chained. The terms "primary" and "secondary" are used here to differentiate between the light-detecting first layer of "on"-detectors, and the second layer of opponent process detectors fed by the first layer. In a multi-resolution structure similar to a pyramidal filter bank analysis, the illumination-resolutions of the different scales of detectors could stack, permitting much finer resolution than would be available from any individual detector. Such detailed luminance-resolution would be dependent only upon adjacent pixel-to-pixel calibration accuracy, but would not require absolute accuracy.

The second idea relates to detector self-calibration, referred to above. For the luminance case (extendable to color channels and any other form of signal, as would be obvious to one skilled in the art), if a row of several detectors "sees" a real spatially varying, temporally stable luminance, and the array is oscillated under the source such that each pixel in the array gets to see each source point, then the systematic relative differences between the detectors can be determined, and the pixels calibrated against their neighbors by adjustable weightings to the opponent layer. If the outer perimeter of primary detectors is exposed to black as a reference, then all pixels can be corrected to an absolute level. If the outer perimeter is shown the average illumination, then the image can be "color-corrected" as in human vision. If one takes a color photograph in fluorescent light, human faces will appear greenish, although to the human eye faces seem normal. This is a visual phenomenon called color-constancy, which could be imitated by this system.

The third idea is that of using the differential image provided by the center-surround structures, and determining position accurately from the exact time of crossing (or phase) of a contrast edge over such a

detector. This can give arbitrarily high positional accuracy independent of the shape or size of the detectors (pixels) themselves, leaving open the possibility of having large-area pixels with good light gathering capability, simultaneously with high spatial resolution of edges extracted from the time domain. To visualize how the detector size can be made irrelevant, one can consider a circular optical detector divided into four detecting quadrants, and illuminated with a spot of light falling entirely on the detector area. If the four quadrants can be compared in pairs (determining ratios of energies falling on pairs of quadrants), then the position of the centroid of the illuminating spot can be known accurately, limited by the accuracy of the comparator, and not by the sizes of the four quadrants, which could be arbitrarily large. Given the self calibration capability of adjacent pixels of a detector used in this manner, timing of edge crossings can give substantially better spatial resolution than pixel dimension alone would provide, thereby reducing required detector complexity.

The fourth idea addresses spatial isolation of textures, edges and points. For a particular orientation of oscillation of the image upon the detector, crossings over center/surrounds by various contrast-amplitudes representing a single slice of texture in that direction will yield a specific frequency spectrum of edge-crossing bursts, dependent upon the spacings and contrast-amplitudes within the texture slice, and the frequency (or frequencies in the case of a chirp) of the generated oscillation. It should be noted that since the contrast contours of the real-world texture can be determined (by oscillation) to higher spatial accuracy at the detector than detector size alone would indicate, more accurate high-frequency information is available than would be present from analysis of a standard bitmap of the same image. Using more than one direction of oscillation, regions with similar texture slice spectra in more than one direction can be grouped. This is another plausible purpose of micro-saccades.

By oscillating the image in the scale domain (zooming in and out), purely radial velocities of moving objects headed towards (or away from) the detector can be simply determined. Thus, using two or more detector arrays, accurate real-time collision-avoidance (or pursuit) systems could be implemented. In the case of robotics, to catch a falling or thrown object (currently impossible using raster-based cameras), adjusting the oscillations (using chirps) on the event-driven, real-time detectors to attain a purely radial resulting velocity (at one point in the chirp) with respect to the catcher, and noting these data, would give sufficient information for object interception.

<u>A more detailed description of the invention follows:</u>

A two-dimensional image detector array has an image focused upon it, which is analyzed in an opponent center/surround structure similar to processes in the human retina (see Figures 1 and 2). The image is thus a differential image. Any luminance edge within the image crossing such an opponent center/surround detector (comprising multiple imaging pixels connected by logic circuits) will trigger a motion detection event (impulse) at the instant that the centroid of an edge passes over the center of such a structure, along with a decaying level representing the instantaneous contrast between the center and its surround. In a desired embodiment of such a detector, the circularly symmetric center/surrounds are composed of various orientations of rows of pixels connected as linear center/surrounds, contributing separately to the concentric circular opponent structure (see Figure 3), similar to ocular dominance column mechanisms in human vision. This would yield orientations of edges and directions of motions. The attainable timing accuracy of the centroid of a spot crossing over an opponent center/surround detector will determine the achievable "resolution" of the edge, not the array's pixel size, or the center, or surround sizes of opponent center/surround structures. If every edge that crosses such a detector is subsequently displayed with the exact relative timing of its crossing, then the perceived spatial resolution would be much higher than that viewed on a raster scan display of the same physical resolution. Alternatively, the implied finer path knowledge of the moving edge could be displayed as a temporal sequence of finer pixel crossing events on a higher-resolution display. This extra temporal knowledge at edges would yield higher spatial resolution information than would be indicated purely by the detector's

pixel size. This permits use of larger pixels at the detector, integrating more signal, yielding higher signal-to-noise ratio (SNR).

Since the image is viewed as a differential luminance map at the opponent layer of the detector system, if the imaging array border is covered by reference-black or preferably a known neutral density wedge fading to reference-black, then all luminance (and similarly, colors) can be coded as spatially accumulated differences from the surrounding black reference across several center/surround detectors, or several scales of center/surround detector. (This reference becomes crucial in the eye, where there is no evidence of a low-pass version of the image, as is available in a wavelet pyramid, to calibrate color and luminance by.) A graded wedge would be useful in allowing a gradual brightening near the edge of the field of view of the detector toward the center, even in a uniformly bright field, permitting adjacent opponent center/surrounds near the edge to report differences kept within the range of each of the individual center/surrounds (see Figure 4). In a multi-resolution structure, such as a wavelet pyramid, this would permit a higher dynamic range than any individual opponent center/surround cell is capable of, due to the successive approximation nature of the increasingly finer scales. The just-noticeable-difference threshold could be spatially (and temporally) varied locally, to reflect the characteristics of human vision, or to reduce information content in low light, high contrast, or high texture, (possibly to reduce computational load on any subsequent intelligent vision system. If the outer perimeter of primary detectors is exposed to black as a reference, then all pixels can be corrected to an absolute level. If the outer perimeter is shown the average illumination, then the image can be "color-corrected" as in human vision. If one takes a color photograph in fluorescent light, human faces will appear greenish, although to the human eye, faces would seem normal under the same lighting. This phenomenon, called color-constancy, permits faster recognition of objects and people than would be possible with radical color shifts dependent upon available illumination such as sunlight, moonlight, or firelight. Ignoring texture, most slowly spatially varying luminances in naturally illuminated scenes can be described by JND contours (JND's being only one bit deep by definition), similar to isobars in on weather maps. Therefore, the dynamic range of each individual pixel does not have to be great in a multi-resolution opponent center/surround structure. This is especially true if each pixel's sensitivity is logarithmic. (To the author, these would seem to be the underlying structures in biological vision, where the total dynamic range is huge, while simultaneously, the sensitivity to tiny adjacent differences is high, yet individual neural detector "bit-depths" must be too small to accurately cover the entire dynamic range of the visual system.)

If dual polarity pairs of opponent center/surrounds are used, covering the same spatial location (again reflecting similar human retinal structures), then an edge traversing an on-center/off-surround will cause a positive impulse, while the same event simultaneously detected in an off-center/on-surround will be indicated by a negative impulse. If on-centers poll a small population of pixels in a small spatial cluster, and off-centers poll a different small population of pixels spatially centered in the same cluster, with surrounds sampled analogously from an annulus of pixels around the central cluster, then a single noise event at a given instant in any pixel contributing to a center (or its surround) will be unlikely to correlate to any simultaneous single pixel noise event contributing to a center (or surround) of the opposite polarity at the same location (see Figure 5). Such a dual detection process would cancel out most detector noise events due to the lack of simultaneous response from the spatially coincident, opposite-sign opponent center/surround. It would also confer immunity to DC level changes analogously to a differential amplifier.

By adding a one- or two-dimensional pre-determined spatial oscillation to the image (or the detector array) intermittently, real edge crossing events (due to static edges in the image crossing these detectors solely because of the relative oscillation) will possess the spatio-temporal motion signature of the induced oscillation. These events can be filtered, since the oscillation is known, or can be determined precisely by a reference illumination spot aimed at a sample point within the detector, feeding back the exact, measured nature of the oscillation. Detector noise will have an extremely small probability of passing

through such a filter. Fast, time-varying image noise can also be removed by such a filter if desired. Also, luminance variations, which are seen to be a systematic function of pixel location on the detector via use of the oscillation, can be nulled, thus calibrating all of the primary detectors relative to their neighbors, and ultimately to the largest opponent surround value, which could be black, the average of the ambient illumination, or any other desired reference.. The author suspects this to be another purpose of micro-saccadic motions seen in the human eye.

The various oscillations could be used to remove unwanted components from an image in real-time based upon their common velocity (and possibly other simultaneously common features such as color). For example, driving in heavy falling snow, the human eye is handicapped by masking noise (the snow is spatially noisy), and masking luminance (the whiteness lowers contrast of the background). The background is quite often "visible" between the snowflakes, but hard to see purely because of masking effects. Since the snow has a common velocity in the temporal domain, those pixels reporting flakes could be suppressed in real-time, and the background contrast brought back to normal on the driver's display.

In between oscillations, any real motion of an edge along a specific orientation of several center/surround detectors will cause a series of centroid crossover events, which may be directly interpreted from the detector as a velocity, in much the same way that a child dragging a stick along a picket fence would indicate speed information by the frequency of the rattling stick along the picket fence. (Since every center/surround crossing event is reported by this detector, there is no missing motion data, and velocity and motion vector fields can be reported exactly.) If the opponent detectors indicating the latest positions of static edges found during the oscillation phase were suppressed between oscillations, the system could be made very sensitive to anything possessing real motion, since only these detectors would be permitted to report motion during the periods between oscillations. This is automatically accomplished in vision because of the time constant of the cones dimming static parts of the image between micro-saccades. In any intelligent vision application, the isolation of known moving targets and the identification of their velocities, are critical time-consuming tasks, which this detector structure could perform efficiently and accurately with little processor burden.

A "chirp" could be put to particular use as the induced oscillation, since any object in motion across the field of view of the imaging array during the chirp would yield a Doppler-shift with higher SNR, giving potentially more accurate velocity information. A camera based upon this concept would only update JND moving edges, and do so instantaneously as they move, since static edges would still be "remembered" by the opponent center/surround detectors, and the matching receiver would possess the information about previous static edges. These detectors could be designed with any reasonable "memory" time constant. Such a system would have access to free, accurate, real-time velocity data, eliminating the need to calculate inaccurate motion vector fields from sequential, time-blurred (possibly interlaced) frames. Currently, motion vector fields are computationally expensive, and require a latency of at least one image frame delay plus the computation time, rendering them unusable in high-speed real-time motion applications. (Information about velocities determined from the differences between frames could be seriously outdated by time it is available, even if the information were accurate for the period being calculated.)

As stated previously, the oscillation or chirp may be used in different orientations in the image to extract multiple directions of one-dimensional, high-resolution texture spectra over different lengths of opponent center/surround detectors at a given coarser scale for rough identification of similar texture areas. The coarse scan can group textures using a crude spectral comparison for purposes of locating approximate borders between different textures. Then, only along these approximate borders, a higher-resolution texture difference scan can be made using more sensitive criteria. This process would not involve massive iterations, or computationally prohibitive filtering, while permitting spatially accurate texture

edges to be mapped. These texture spectra would be derived from the pulse trains given by the chained opponent center/surround detectors having the textures oscillated over them, and would reflect their spatial resolution. By using a multi-resolution center/surround structure, starting with the coarser scale, regions may be grouped by similarity, then at a finer scale near texture-difference borders previously derived from the coarse scale examination, exact texture borders may be determined, using less on-chip processing. Thus, similar texture regions may be efficiently isolated, and, given a suitable look-up table, identification of textures is of course possible.

It is surmised by the author that spatial processing of texture occurs first at the retina (possibly to limit processing at, and data transmission to the visual cortex), but that there is insufficient neural structure or bandwidth to handle the texture directly in the retinal image-space. Micro-saccades would serve the purpose, by transforming the two-dimensional image structure into a few one-dimensional structures (neural pulse trains) "mixed-down" to the lower frequencies handled by neurons. Texture could be differentiated from edges and isolated points as follows: given a particular size of center/surround region, the oscillation of a texture in various directions would yield orientation-dependent noise spectra, while oscillating across a simple edge would give a much narrower, shaped band response, dependent upon the orientation of oscillation, and an isolated point would give back a narrow spike for any orientation of oscillation (See Figure 6).

Oscillation of the image in the scale domain gives access to radial velocity information regarding moving objects headed towards (or away from) the detector. Therefore, using two or more detector arrays, accurate real-time collision-avoidance (or pursuit) systems could be implemented in inexpensive hardware. The application for civil aviation and for blind-spot coverage in automotive applications is enormous. As mentioned above, pseudo-human robotics capabilities are made possible by this architecture.

A stereoscopic pair (or more) of these proposed detector arrays could be made to provide information such as safe landing locations for pilots, based upon texture analysis, or instantaneous trajectories for path prediction, from accurate stereoscopic velocities, useful for navigating (e.g. robotic arms, vehicles, aircraft, etc.) in three dimensional space for interception/avoidance and capture tasks, since the proposed array would report every detector crossing event in real-time with accurate positions and velocities. This is a capability not available in vision systems having to process raster scanned images, and extract inaccurate and delayed motion vector fields. Indeed, a delay of one video frame alone, in any real-time system makes navigation and interception of moving objects very difficult for anything but constant velocity or low velocity applications. Three detector arrays that are monitored jointly, for example, will provide important additional information in some circumstances that is not obtainable by using only two detector arrays.

References:

C Language Algorithms for Digital Signal Processing, P.M. Embree, B. Kimble, Prentice Hall, 1991
Fundamentals of Nonlinear Digital Filtering, J. Astola, P. Kuosmanen, CRC Press, 1997
Digital Image Processing, R.C. Gonzalez, R.E. Woods, Addison Wesley, 1993
Wavelets Algorithms & Applications, Y. Meyer R.D. Ryan, SIAM, 1993
How the Mind Works, S. Pinker, Norton, 1997
Visual Science and Engineering Models and Applications, D.H. Kelly Editor, Marcel Dekker, 1994
Vision, D. Marr, W.H. Freeman and Company, 1982
Eye, Brain, and Vision, D.H. Hubel, Scientific American Library, 1995
Computational and Psychophysical Mechanisms of Visual Coding, M. Jenkin and L. Harris Editors, Cambridge, 1997
Computer Vision Algorithms in Image Algebra, G.X. Ritter, J.N. Wilson, CRC Press, 1996
Colour: Its Measurement, Computation and Application, G.J. Chamberlin, D.G. Chamberlin, Heyden, 1980
Color Science, G. Wyszecki, W.S. Stiles, John Wiley & Sons, 1982
The Image Processing Handbook Second Edition, J.C. Russ, CRC Press, 1995
Digital Video Processing, A.M. Tekalp, Prentice Hall, 1995
Two-Dimensional Signal and Image Processing, J.S. Lim, Prentice Hall, 1990
Texture & Modeling, A Procedural Approach Second Edition, D.S. Ebert et al, Academic Press, 1994
Reference Data for Radio Engineers, E. Baguley et al, Howard W. Sams & Co., 1979
Electronics One-Seven, H. Mileaf Editor In Chief, Hayden, 1967

Prior Art:

No directly related prior art is known to exist.

Claims:

I Claim:

1. A method of extracting information not available in a static image or a raster-scanned image sequence, by oscillating an image (spatially and/or in scale) and a suitable-architecture detector array relative to each other in a controlled fashion, (whether the image is optical, acoustic, is composed of radio waves, x-rays, pressure, or any other medium, and whether by moving the detector or by altering the image path), and where the oscillation is a sinusoid, a chirp (swept frequency), or any other suitable time varying motion in one or more dimensions, or in scale.

2. The method according to Claim 1, where the extracted information is phase information of stationary objects or edges crossing secondary detector boundaries due to detector array oscillation, giving increased spatial accuracy not available from primary detector size alone where said phase information is generated in an event driven manner, not in raster or frame sequence.

3. The method according to Claim 1, where the extracted information is instantaneous velocity information of moving objects or edges crossing secondary detector elements during induced oscillation, by Doppler shift detection of opponent detector crossing frequency, or by direct timing of detector crossings either during the detector's motionless periods, or during the induced oscillation, compensating for said detector oscillation.

4. The method according to Claim 1, where the extracted information is the pure radial approach or departure velocity of objects requiring collision-avoidance or interception, and/or the information that they are, or are not candidates for collision.

5. The method according to Claim 1, employing two (and possibly more) such detector arrays, where the extracted information is the trajectory and velocities required for interception of a given target at a given point in space-time.

6. The method according to Claim 1, where the extracted information is of true objects or edges versus space-time varying noise events at each detector element, for the purpose of increasing the signal to noise ratio of the system.

7. The method according to Claim 1, where the extracted information is the instantaneous, systematic detector-to-detector sensitivity variations, for the purposes of relative or absolute calibration of the array elements.

8. The method according to Claim 1, where the extracted information is a measure of the local contrast and/or local velocity, useful for spatially adapting the sensitivity threshold of the detector array in real-time, for any of the purposes of matching human sensitivity, gain control, dynamic range extension, entropy control, background or texture suppression, or any other application capable of using such real-time information.

9. The method according to Claim 1, where the extracted information is the identification of moving objects or edges, by suppression of detection of stationary objects or edges (where suppression of static objects may be achieved either actively, or by time-constant governed decay of static images during still-periods of the detector array), to reduce the processor burden in tasks requiring tracking, identification, interception, or any other task which would be slowed or complicated by the lack of such information.

10. The method according to Claim 1, where the extracted information is the various spatial luminance/contrast spectra sampled at various orientations and/or scales, of different textures exposed to the detector array, derived from the relative motions in each orientation, of the textures upon the detector array, whether that information is used for texture identification, segmentation, or any other purpose.

11. The method according to Claim 1, where the extracted information is the various spatial luminance/contrast spectra by orientation and scale, of different textures, edges, and/or isolated points, exposed to the detector array, derived from the relative motions in each orientation, of the features upon the detector array, whether that information is used for feature identification, segmentation, or any other purpose.

12. The method according to Claim 1, employing two (and possibly more) such detector arrays, to accurately identify real and predicted trajectories of moving objects or edges in three-space in real-time, from the differing instantaneous apparent velocities on the separate detector arrays, for the purposes of tracking, manipulation, interception, avoidance, or any other synthetic vision, vehicle guidance or robotics task.

13. The method according to Claim 1, where a neutral density wedge fading to reference-black towards the edge of the array (where "black" refers to non-measurable or suitably low illumination by whatever light, sound, pressure, or electromagnetic energy is being measured by the array), surrounds the perimeter of the detector array, for the purpose of limiting or preventing opponent-detector saturation.

14. The method according to Claim 1, where a neutral density wedge fading to reference-black towards the edge of the array (where "black" refers to non-measurable or suitably low illumination by whatever light, sound, pressure, or electromagnetic energy is being measured by the array), surrounds the perimeter of the detector array, for the purpose of extending the total dynamic range of the array, by allowing opponent contrasts to stack

15. The method according to Claim 1, where a neutral density wedge fading to reference-black towards the edge of the array (where "black" refers to non-measurable or suitably low illumination by whatever light, sound, pressure or electromagnetic energy is being measured by the array), surrounds the perimeter of the detector array, for the purpose of permitting black-reference calibration of the entire detector array in the absence of any stored reference data, or any low-pass version of the viewed image.

16. The method according to Claim 1, where a neutral density wedge fading in ambient illumination, (where the illumination represents light, sound, pressure or electromagnetic energy is being measured by the array), surrounds the perimeter of the detector array, for the purpose of permitting ambient-normalized reference calibration of the entire detector array in the absence of any stored reference data, or any low-pass version of the viewed image.

17. The method according to Claims 1, 6 and 7, wherein detector noise, cross-talk, and image noise are sufficiently lowered, and pixel-to-pixel calibration is sufficiently accurate, to permit vision in substantially reduced illumination, in applications such as, but not limited to, night-vision (where illumination is risky, low, or otherwise not available), radiology (where illumination implies increased radiation exposure), and radar (where illumination implies loss of stealth and/or higher power levels for the same SNR or range).

18. A system for retrieving data from an image using a detector array, said system comprising a detector array, with means to spatially oscillate said detector array relative to said image in a controlled manner and means to retrieve data from said image through said detector array.

19. A system as claimed in Claim 18 wherein said detector array has opponent centre/surround components.

20. A system as claimed in Claim 19 wherein said detector array has a first layer of detector elements and a second layer of opponent centre/surround components, said second layer receiving positive or negative input from one or more elements in said first layer.

21. A system as claimed in Claim 20 wherein said first layer of detector elements is sensitive to at least one of luminance, electromagnetic energy, pressure, sound or any other measurable quantity, with means to feed said secondary layer from one or more detector elements of said first layer.

22. A system as claimed in Claim 21 wherein said second layer components are grouped by spatial orientation, said secondary layer components extracting contrast information from the detector elements of the first layer, said second layer components extracting instantaneous spatial object motion and velocity information.

23. A system as claimed in Claim 22 wherein there are two detector arrays positioned in different spatial locations, said detector arrays being jointly monitored to accurately identify and predict trajectories of moving objects or edges in three dimensions in real-time by retrieving the different instantaneous apparent velocities from said detector arrays.

24. A system as claimed in Claim 23 wherein there is a third detector array positioned in a different spatial location from said two detector arrays, said detector arrays being jointly monitored.

25. A system as claimed in Claim 19 wherein there is a neutral density wedge located adjacent to an edge of said detector array and extending around a perimeter thereof, said wedge limiting or preventing opponent-detector saturation.

26. A method for retrieving data from an image using a detector array, said method comprising spatially and/or in scale oscillating said detector array relative to said image in a controlled manner and retrieving data from said image through said detector array.

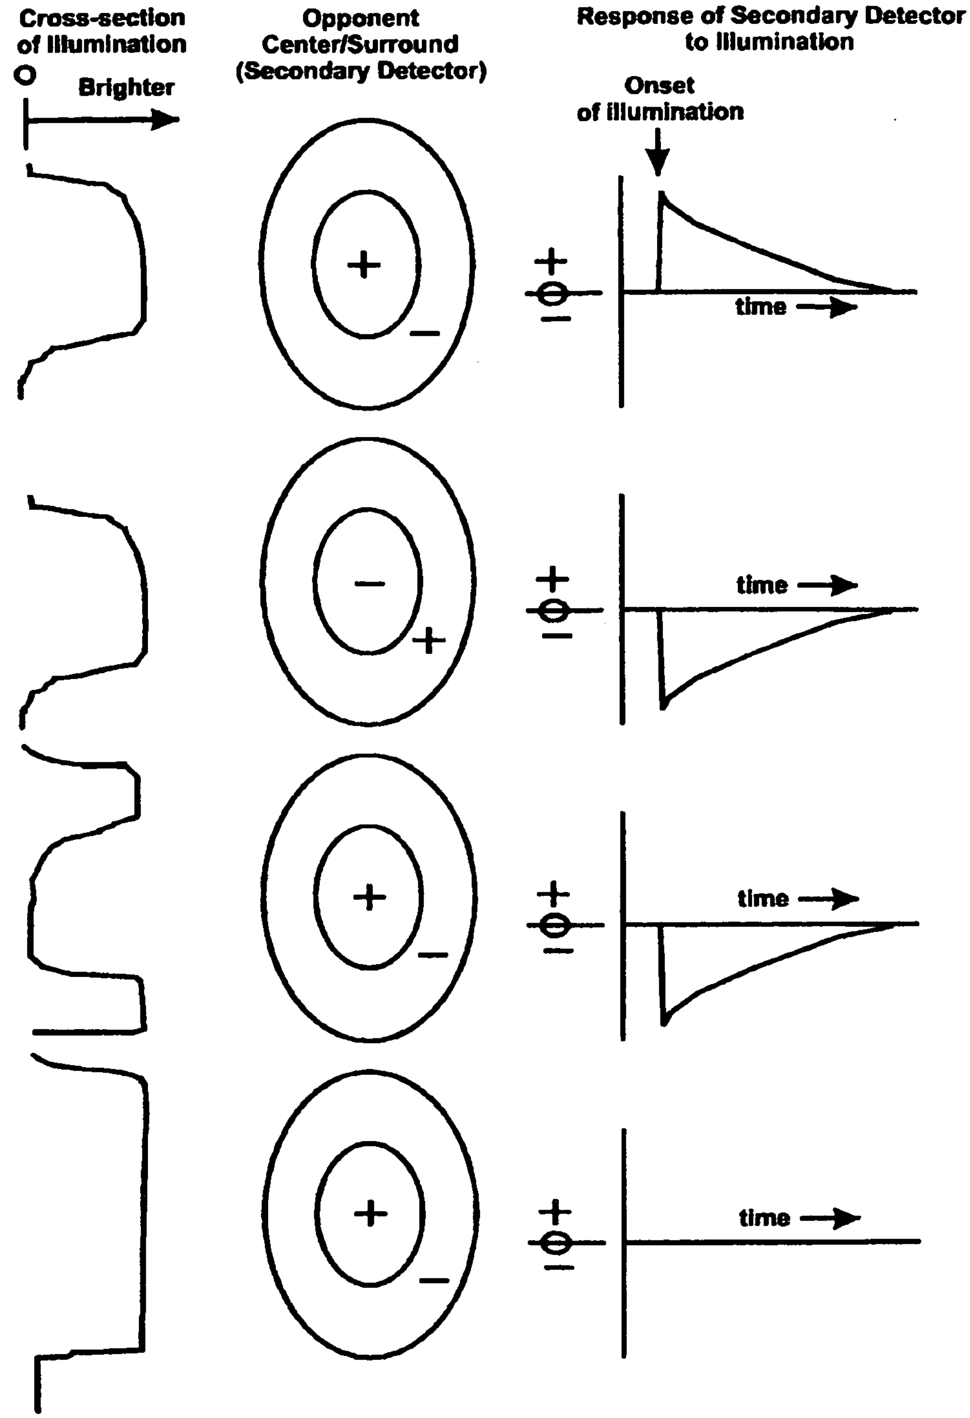

Figure 1

Example Illuminations of Proposed Opponent Center/Surround Detectors and Resulting Outputs.

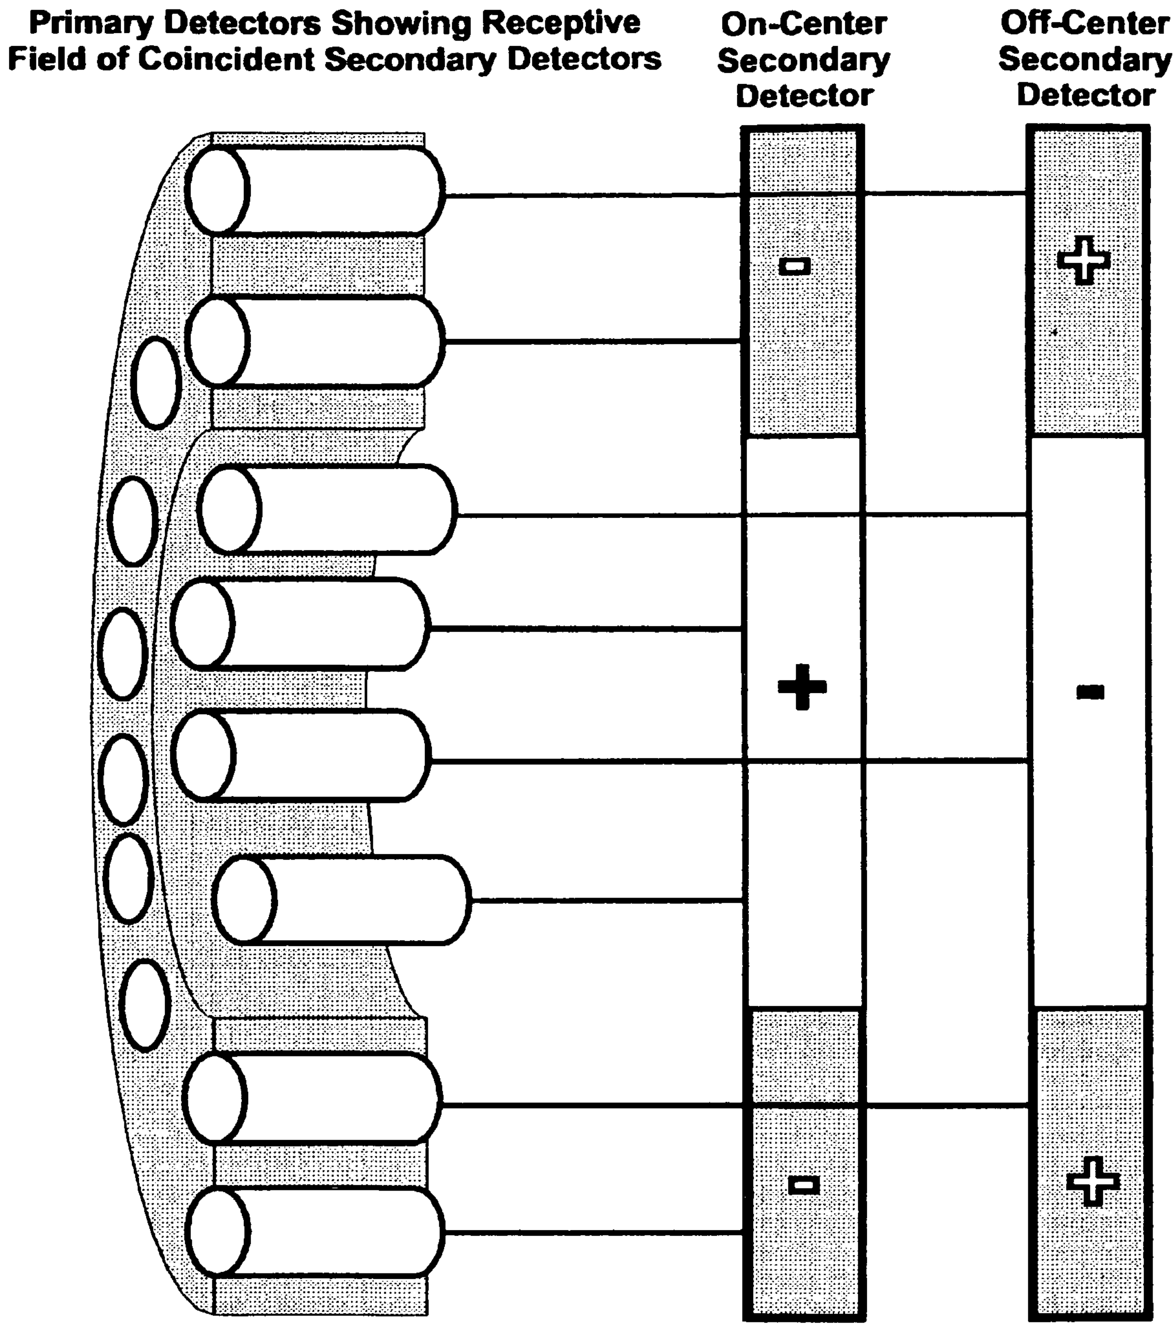

Figure 2

Concentric Regions of Primary Detector Sampled by Proposed Opposite-Polarity Secondary Detectors.

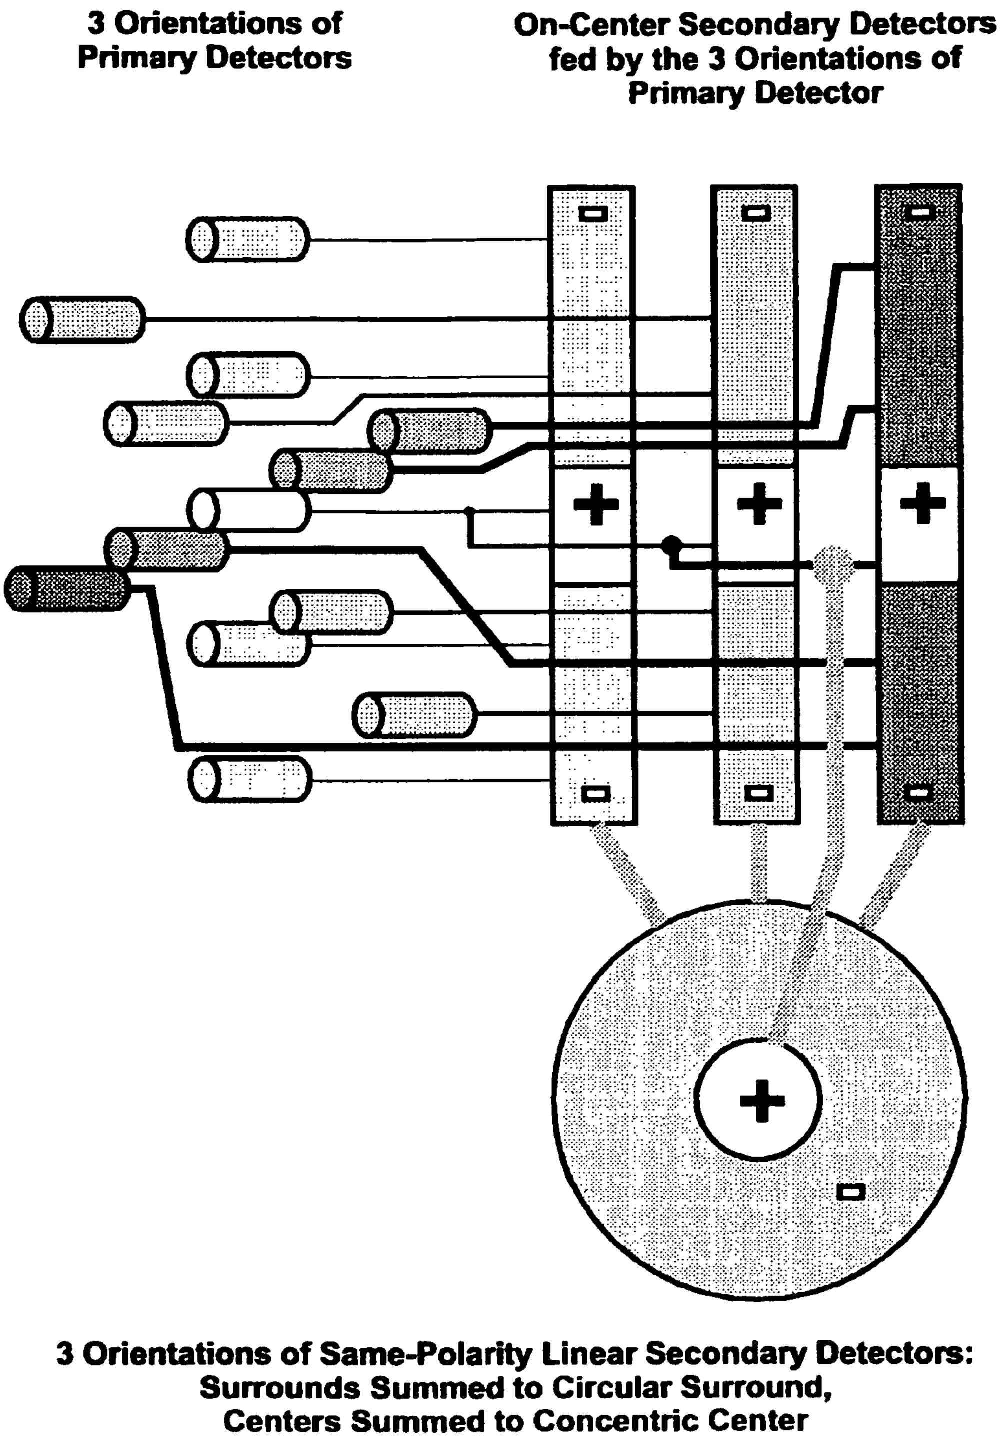

Figure 3

Proposed Linear Opponent Center/Surrounds Contributing to Circular Concentric Opponent Detector.

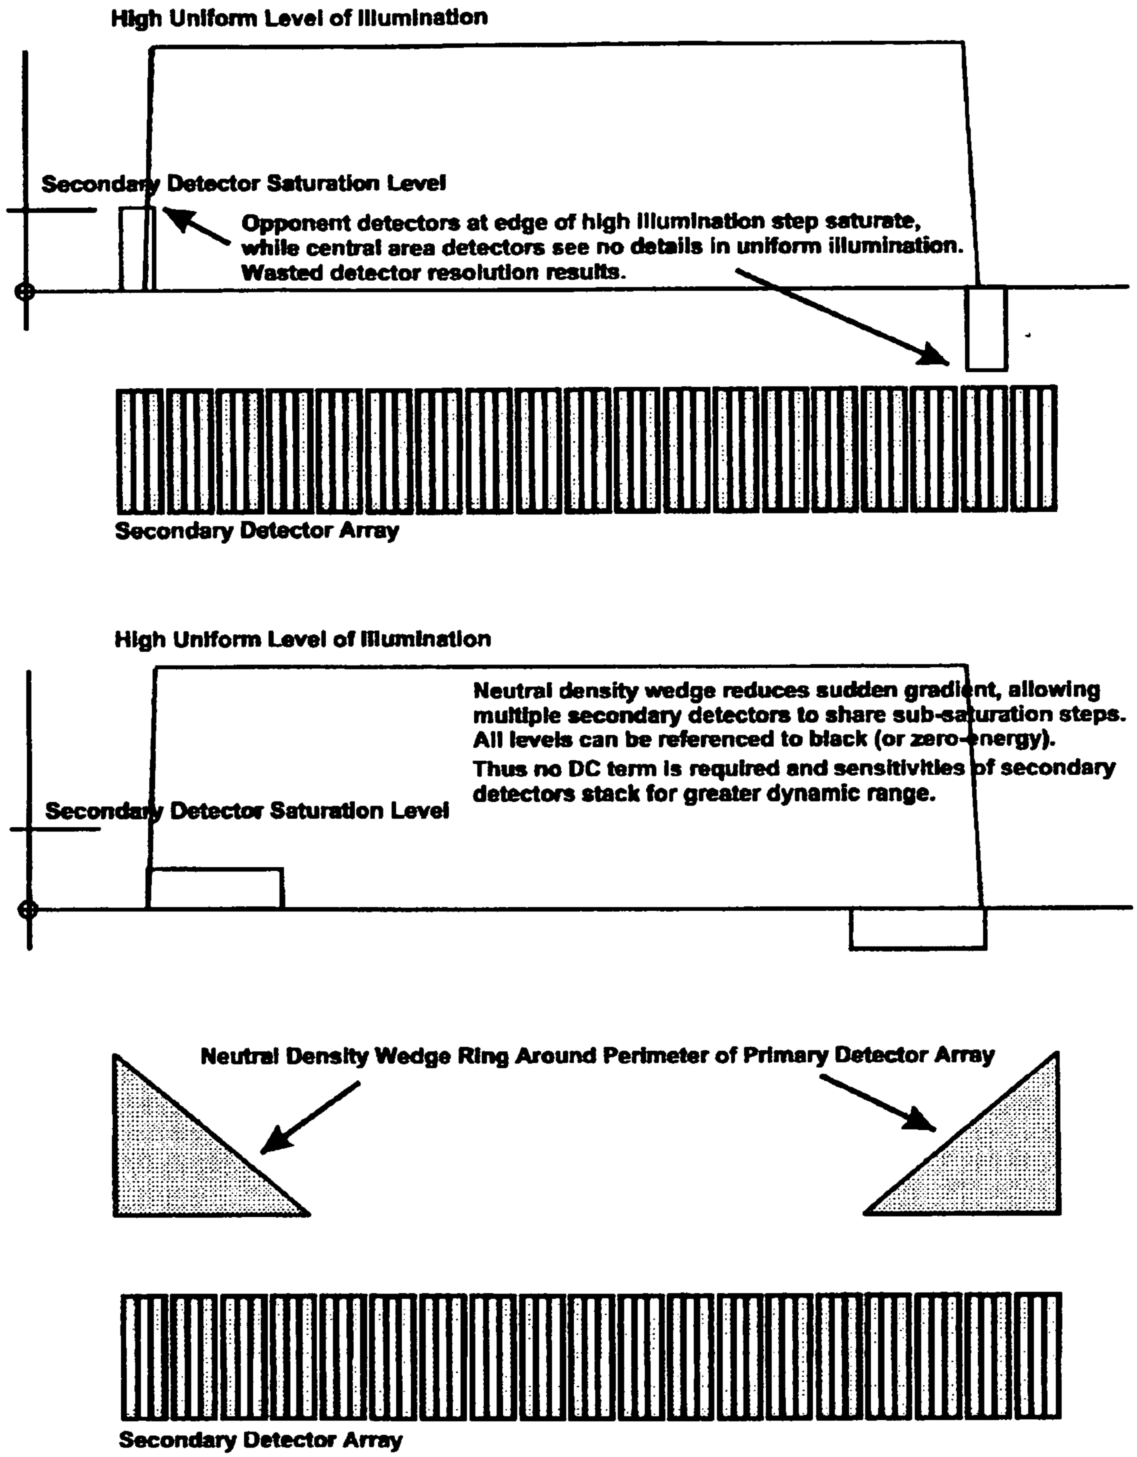

Figure 4

Neutral-Density Perimeter Wedge to Prevent Saturation and Increase Dynamic Range of Sensor Array.

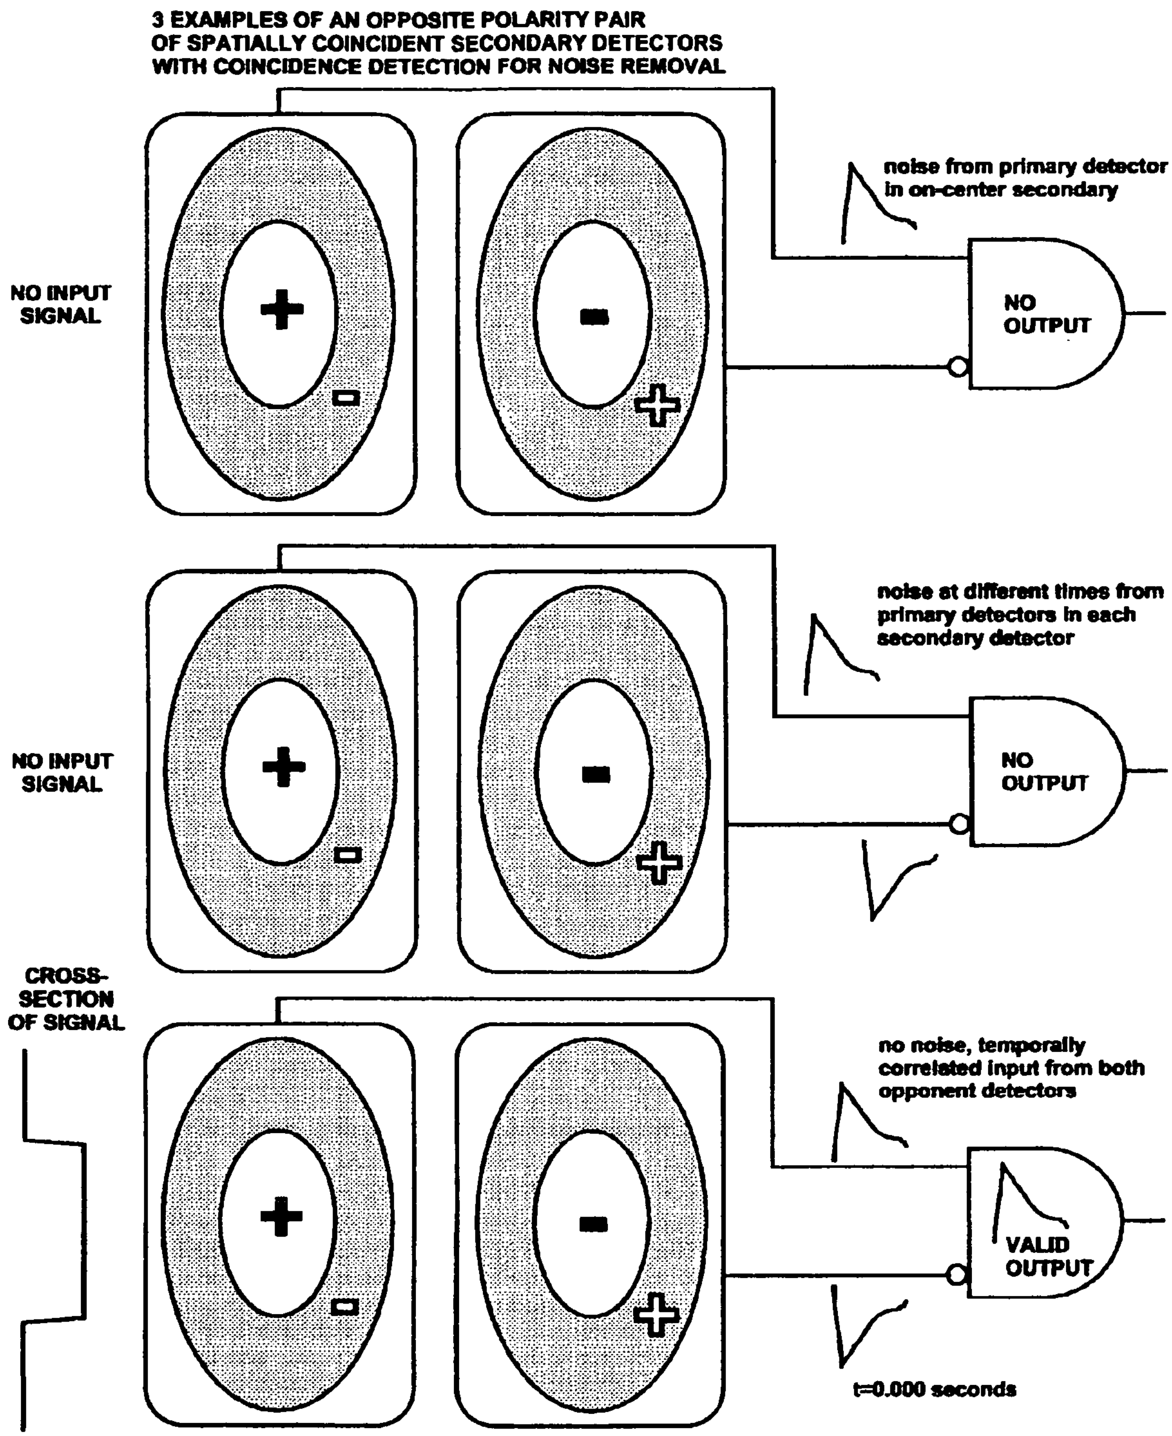

Figure 5

Spatially-Coincident Opposite-Polarity Secondary Detectors for Cancellation of Detector Noise.

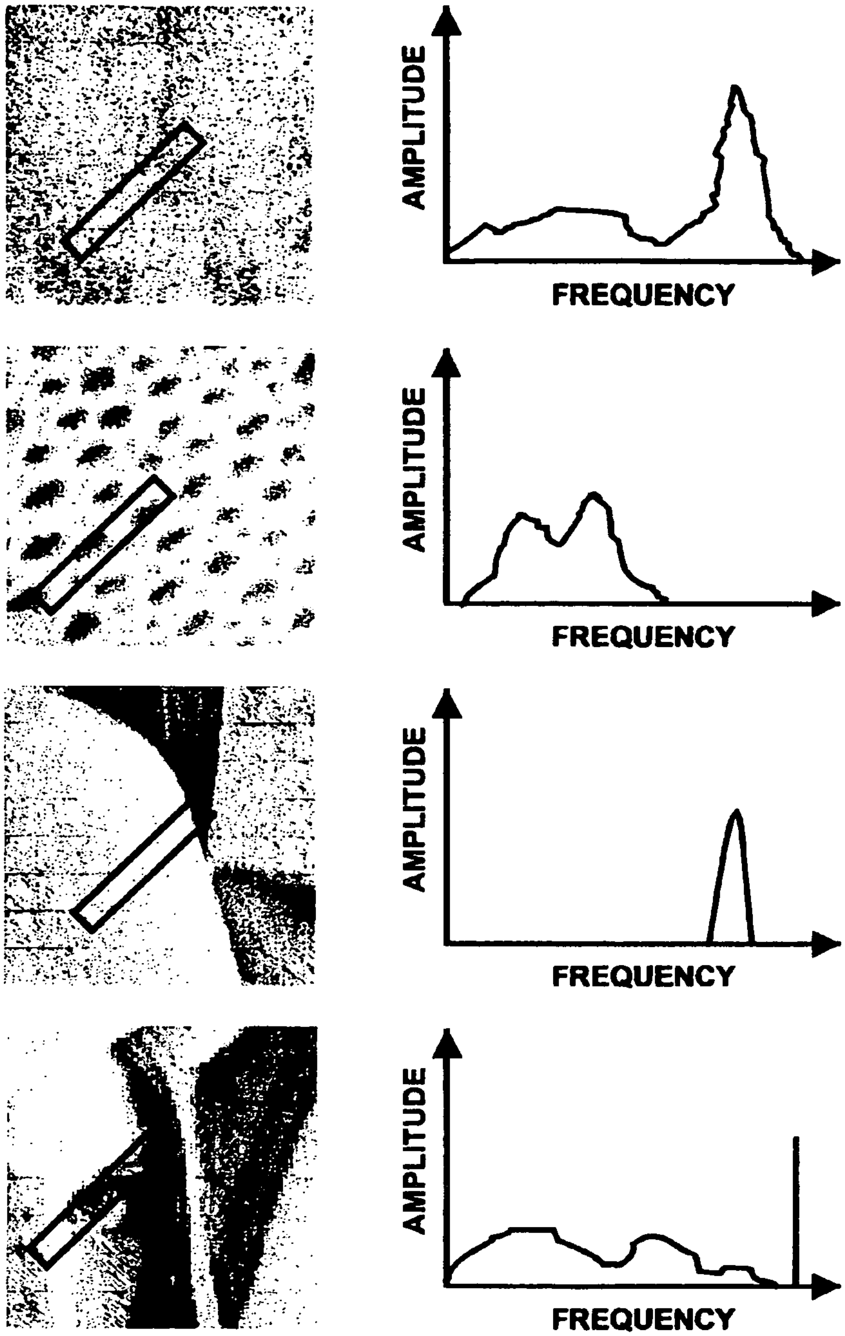

Figure 6

Two Textures, an Edge, and a Pixel Highlight, Oscillated at One Orientation, with Possible Spectra. Note: Actual scenes being sampled by the oscillation will be differential images due to the secondary opponent center/surround detectors, not the direct luminance scenes depicted above. Spectra taken in the differential image space may be interpreted directly, or converted back into the image domain.